Frequency domain reduced order model of aligned-spin effective-one-body waveforms with generic mass-ratios and spins

Michael Pürrer[1,2]

[1]School of Physics and Astronomy, Cardiff University, Queens Building, CF24 3AA, Cardiff, United Kingdom
[2]Max Planck Institute for Gravitational Physics (Albert Einstein Institute), Am Mühlenberg 1, D-14476 Potsdam-Golm, Germany

I provide a frequency domain reduced order model (ROM) for the aligned-spin effective-one-body (EOB) model "SEOBNRv2" for data analysis with second and third generation ground based gravitational wave (GW) detectors. SEOBNRv2 models the dominant mode of the GWs emitted by the coalescence of black hole (BH) binaries. The large physical parameter space (dimensionless spins $-1 \leq \chi_i \leq 0.99$ and symmetric mass-ratios $0.01 \leq \eta \leq 0.25$) requires sophisticated reduced order modeling techniques, including patching in the parameter space and in frequency. I find that the time window over which the inspiral-plunge and the merger-ringdown waveform in SEOBNRv2 are connected is discontinuous when the spin of the deformed Kerr BH $\chi = 0.8$ or the symmetric mass-ratio $\eta \sim 0.083$. This discontinuity increases resolution requirements for the ROM. The ROM can be used for compact binary systems with total masses of $2M_\odot$ or higher for the advanced LIGO (aLIGO) design sensitivity and a 10 Hz lower cutoff frequency. The ROM has a worst mismatch against SEOBNRv2 of $\sim 1\%$, but in general mismatches are better than $\sim 0.1\%$. The ROM is crucial for key data analysis applications for compact binaries, such as GW searches and parameter estimation carried out within the LIGO Scientific Collaboration (LSC). **LIGO-P1500175**

PACS numbers: 04.25.Dg, 04.25.Nx, 04.30.Db, 04.30.Tv

arXiv:1512.02248v1 [gr-qc] 7 Dec 2015

I. INTRODUCTION

The Advanced LIGO (aLIGO) gravitational wave (GW) detectors [1] have recently started their first observing run and Advanced Virgo [2] is expected to start taking data in 2016. The first direct detection of gravitational waves is expected to be made before the end of this decade when these detectors should reach their design sensitivity [3]. The Einstein Telescope is a planned third generation detector and will be an order of magnitude more sensitive that second generation detectors once it becomes operational after 2025 [4].

Coalescences of binary black hole (BH) systems are one of the most promising sources for ground based gravitational wave detectors. Possible GW signals will be buried in the noisy output of the interferometers. Matched filtering can be used to look for signals in GW searches by correlating the detector data against GW templates [5]. Parameter estimation then attempts to extract more accurate information on astrophysical parameters of the binary.

A prerequisite for searches and especially parameter estimation is the availability of fast and accurate models of the GW waveform emitted from BH binaries so as not to miss signals or misrepresent their astrophysical parameters. Typically $10^6 - 10^7$ waveform model evaluations are needed for a PE simulation. Stochastic template banks algorithms [6, 7] for GW searches typically produce banks up to $N \sim 5 \times 10^5$ templates [8–10] in the aligned-spin parameter space. In the construction of these banks many more waveforms need to be generated. The authors of [6] suggest that N^2 evaluations are used to make sure the final bank is effectual. This number of evaluations is not feasible even for the fastest waveform models and the number used in practise is estimated to be $\sim N^{1.5} \approx 10^8$ waveform evaluations [11]. While the construction of the bank is very costly, it can be used as long as the power spectral density (PSD) of the detector noise does not change too much [12].

Inspiral-merger-ringdown models of GWs emitted from BH binaries have been constructed either in the effective-one-body (EOB) approach [13–17] or as phenomenological models [18–21]. In this paper I focus on the "SEOBNRv2" [22] model that describes GWs from binaries with spins aligned with the orbital angular momentum. Compared to its predecessor "SEOBNRv1" [23] it extends the spin range to almost the entire Kerr range.

EOB waveforms are obtained from the integration of complicated systems of ordinary differential equations and are in general much too slow for direct data analysis applications. Reduced order modeling (ROM) can provide fast and accurate surrogates for such GW models which are crucial for GW searches and parameter estimation. I describe here a ROM that speeds up the generation of SEOBNRv2 waveforms by a factor of thousands. This allows parameter estimation studies to be carried out on the order of a day instead of a year. ROM approaches are based either on the singular value decomposition (SVD) and different interpolation techniques [24–27] or on the greedy basis method and empirical interpolation [28–31].

In this paper I extend the reduced order modeling techniques presented in [24] to SEOBNRv2 waveforms and provide a ROM that covers the *entire* parameter space domain on which SEOBNRv2 is defined and a frequency range sufficient for *all* expected compact binary sources including binary BHs, neutron star - black hole, and binary neutron star systems, while neglecting equation of state effects. A unified ROM covering the mass space for all compact binary coalescences is of great utility for data analysis, even though SEOBNRv2 models GWs emitted from BH binaries and does not include tidal effects or describe tidal disruption. In principle, such effects could be added post hoc as modifications of the amplitude and phasing of SEOBNRv2 [32, 33]. The ROM provides speedups of up to several thousands over the SEOBNRv2 model while being accurate to better than 1% in the

mismatch to keep ROM errors below the calibration accuracy of SEOBNRv2 against NR waveforms [22].

The approach to build frequency domain reduced order (or surrogate) models presented in [24] contains the following important steps: EOB input waveforms are generated on a tensor product grid over the domain of intrinsic parameters (mass-ratio, aligned spins). The waveforms are tapered and Fourier transformed and split into their amplitude and phase parts. These functions can be accurately represented on sparse frequency grids with only O(100) nodes. The total mass does not need to be a parameter when working with geometric frequency Mf. I build separate reduced bases for the amplitude and phase with the *singular value decomposition* (SVD). To construct a waveform model the expansion coefficients of the amplitude and phase as a function of physical parameters are needed. I compute these coefficients from projecting the input amplitudes and phases onto their reduced bases and interpolate the coefficients by tensor product cubic splines. Finally I assemble the ROM from its constituent parts and compute the strain from the amplitude and phase interpolants.

In this paper I extend these techniques for the more challenging SEOBNRv2 case. In particular, the aligned spin range is extended to go close to extremal aligned spins, $\chi_i \in [-1, 0.99]$. As will be discussed later this extended region turns out to be more challenging to model than expected due to peculiarities of the SEOBNRv2 model and it serves as a nice illustratation of how ROM methods can help find problems in waveform models. Crucially, to obtain a faithful model over a large parameter space I employ *domain decomposition techniques* both in the mass-ratio and spins and in frequency to increase resolution at high spins and for high frequencies. Domain decomposition methods and "hp-greedy algorithms" for adaptive sampling for large problems have been discussed in Refs. [28, 34–37].

The organization of this paper is as follows: In section II I summarize the SEOBNRv2 waveform model. In section III I describe new reduced order modeling techniques introduced in this paper and discuss in section IV how the SEOBNRv2 ROM is constructed. Results on speedup and accuracy of the ROM are discussed in section V. Conclusions are given in section VI.

II. THE SEOBNRV2 MODEL

The gravitational wave signal emitted from coalescing binary black holes can be characterized by three different stages. During the long inspiral the orbital velocities are small compared to the speed of light and this stage is well described by post-Newtonian theory [38]. Driven by gravitational radiation reaction the companions approach each other more and more and eventually merge. In that regime, non-perturbative solutions of Einstein's equations are needed which can be computed by numerical relativity simulations [39]. The ringdown of the remnant Kerr black hole and the emission of quasinormal modes can be described by perturbation theory [40].

The effective-one-body formalism introduced by Buonanno and Damour [13–17, 41] combines post-Newtonian theory, resummation techniques and perturbation theory with the help of mapping the two body problem to an effective one body problem. This approach assumes that the comparable-mass behavior is a smooth deformation of the test-particle limit. Three main ingredients of EOB are the conservative two-body dynamics, radiation-reaction, and the gravitational waveform emitted during inspiral, merger and ringdown.

The full EOB waveform is obtained by joining the inspiral-plunge waveform and the merger-rindown waveform over a predetermined interval in time. Over the years the formalism has evolved to incorporate spin orbit and spin-spin interactions, and several high order parameters are calibrated from NR information [42, 43].

In the following I summarize the aligned-spin waveform model "SEOBNRv2" introduced by Taracchini et al. [22] as an extension of an earlier "SEOBNRv1" [23] model that was limited to dimensionless spins smaller than 0.6. SEOBNRv2 is a model for the dominant $l = m = 2$ GW mode emitted from BH binary systems with spins aligned with the orbital angular momentum. The dimensionless component spins $\chi_i = \vec{S}_i \cdot \hat{L}/m_i^2$ cover almost the full Kerr range $-1 \leq \chi_i \leq 0.99$. $\vec{S}_i$ are the spin vectors, $\hat{L}$ the orbital angular momentum unit vector and m_i are the component masses of the BHs. The mass-ratio $q = m_1/m_2 \geq 1$ is allowed to range between equal mass and quite extreme systems at $q = 100$. In terms of the symmetric mass-ratio $\eta := m_1 m_2/(m_1 + m_2)^2$ the model covers approximately $0.01 \leq \eta \leq 0.25$.

The model was calibrated against the $l = m = 2$ mode of a catalog of 38 BH binary NR waveforms produced by the SXS Collaboration [44–47]. In particular, 8 nonspinning NR waveforms with $q = 1, 1.5, 2, 3, 4, 5, 6, 8$, plus 19 spinning NR waveforms were used for calibration. In addition the model was tuned against Teukolsky-based merger-waveforms for $\eta = 0.001$ and several spin values for the central BH. These calibration points are shown in Fig. 1 as a function of the symmetric mass-ratio and a spin parameter

$$\chi := \frac{\chi_{\rm Kerr}}{1-2\eta} = \frac{\chi_1 + \chi_2}{2} + \left(\frac{\chi_1 - \chi_2}{2}\right)\left(\frac{m_1 - m_2}{m_1 + m_2}\right)\left(\frac{1}{1-2\eta}\right), \tag{2.1}$$

where $\chi_{\rm Kerr} = (\vec{S}_1 + \vec{S}_2) \cdot \hat{L}/M^2$ is the dimensionless spin of the deformed Kerr geometry used in the EOB construction.

The model was found to be accurate to a mismatch of 1% against the same set of SXS NR waveforms for total masses in $[20, 200]M_\odot$ and the aLIGO design power spectral density (PSD) [48] with a low frequency cutoff of that depends on the length of each NR waveform (see Fig. 1 in Ref. [22]). A study focused on NSBH systems [49] also found mismatches of 1% or better between SEOBNRv2 and NR waveforms.

A recent study found evidence [19] that while SEOBNRv2 is in general extremely accurate it may not accurately describe the merger-ringdown regime for high aligned spins $\chi_i \gtrsim 0.7$ where it was not calibrated to NR simulations. This is corroborated by a further waveform comparison study [50] with a catalog of new SXS NR waveforms [51].

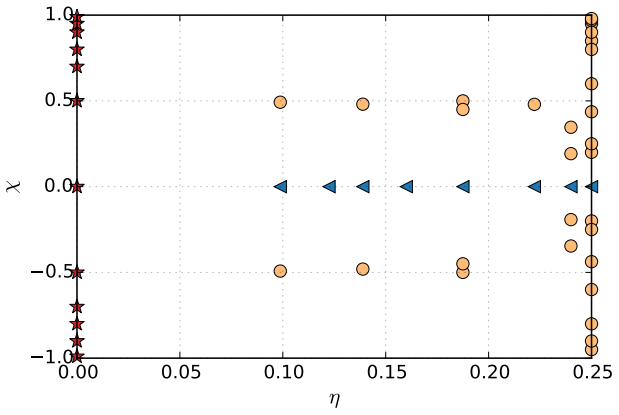

FIG. 1: The 38 SXS nonspinning (blue) and spinning (orange) NR simulations used to calibrate SEOBNRv2 as a function of symmetric mass-ratio η and the spin parameter χ. In addition, the Teukolsky waveforms are shown in red.

III. REDUCED ORDER MODELING TECHNIQUES

In the following I discuss extensions to the reduced order modeling approach given in [24]. In particular, I focus on techniques to achieve increased resolution in some parts of the model parameter space and on an efficient way of handling different resolution requirements over the frequency domain. This allows building accurate models at an economical cost.

As discussed in [24], I build separate bases for the frequency domain amplitude and phasing of the original waveform model using the singular value decomposition. To obtain a ROM a prediction of amplitude and phase coefficients is needed along with the reduced bases. This can be achieved by interpolating the coefficients over the parameter space. Tensor product splines are a robust tool to achieve this, if the density of waveforms in the multi-dimensional regular grid is sufficiently high to capture the variations in the waveform model.

The final product is a surrogate for the GW strain for the $l = m = 2$ mode in the frequency domain [24]

$$\tilde{h}(\lambda; M; f) := A_0(\lambda, M)\, I_f\left[\mathcal{B}_{\mathcal{A}} \cdot I_{\otimes}[\mathcal{M}_{\mathcal{A}}](\lambda)\right] \exp\left\{i\, I_f\left[\mathcal{B}_{\Phi} \cdot I_{\otimes}[\mathcal{M}_{\Phi}](\lambda)\right]\right\}, \quad (3.1)$$

with the intrinsic model parameters $\lambda = \{\eta, \chi_1, \chi_2\}$, the reduced bases $\mathcal{B}$, projection coefficient tensors $\mathcal{M}$, tensor product $I_{\otimes}[]$ and frequency spline interpolation operators $I_f[]$. The two polarizations of the dominant $l = m = 2$ mode of the GW strain are then given by Eq. (6.14) and (6.15) of Ref. [24].

A. Patching in the parameter space

I will show in Sec. V C that resolution requirements for SEOBNRv2 are not uniform over the parameter space. Tensor product interpolation does not allow for local refinement regions, as it is defined as the tensor product of one-dimensional sets of points. A simple improvement can be obtained by decomposing the domain into multiple overlapping patches. This idea is illustrated in Fig. 2. A base patch, covering the whole domain of interest, is complemented by one (or more) smaller refinement patches for regions of higher waveform variability. These regions can be found quite naturally when checking the accuracy of a ROM covering the whole domain.

The combined ROM then switches between ROMs defined on subdomains based on the location where the ROM is to be evaluated. It is prudent to switch between these sub-ROMs not right at the boundary, but somewhat inside the subdomain in question to avoid possible boundary effects. The subdomains chosen for the SEOBNRv2 ROM are described in Sec. IV.

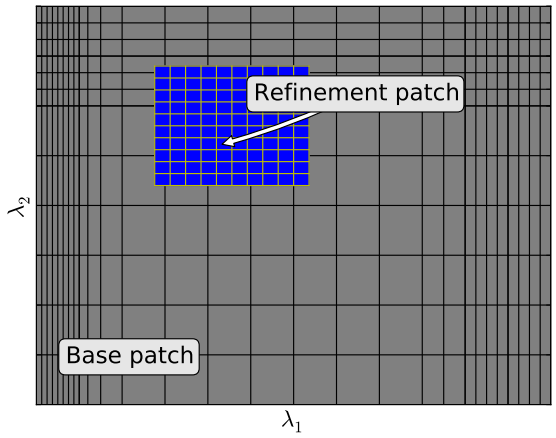

FIG. 2: An example tensor product grid on a base patch covering the entire domain of interest in model parameters λ_1, λ_2 and a refinement patch.

B. Patching in frequency

Resolution requirements also depend on frequency. Due to highly non-linear physics near merger and calibration against NR waveforms it is natural to assume that SEOBNR models will contain more structure at high frequencies, while the behavior at low frequency is expected to be very smooth based on the resummed post-Newtonian information used in the construction of the EOB model. This can indeed be seen by studying the performance of a ROM which covers a broad frequency range at different total masses with an appropriate detector power spectral density. In particular I find that SEOBNRv2 does not smoothly depend on mass-ratio and spins everywhere in the parameter space (see Sec. V C for details).

Since the cost of EOB waveform generation rises steeply as the starting frequency is decreased, the above resolution requirements imply that it is efficient to split the frequency domain into two parts. One ROM can cover the low frequency part with reduced resolution over the parameter space, while another can explore the finer structure present at high frequencies with a high density of waveforms that can be generated in an affordable manner. For the present work these high fre-

quency waveforms are generated at a total mass of $100M_\odot$. In principle, the frequency domain could be partitioned into more than two subdomains, but since each subdomain introduces overhead costs due to interpolation over the parameter space at model evaluation, I restrict myself to two parts.

The input waveforms used to build a ROM using the method outlined in [24] must be defined on a fixed grid in geometric frequency. In particular, there is a fixed geometric starting frequency determined from the lowest total mass of interest, $2M_\odot$ for BNS systems, and from the low frequency cutoff of the detector noise spectrum, which is close to 10 Hz for aLIGO design. For configurations with high mass-ratio this would imply component masses way below $1M_\odot$, a lower bound for the expected mass of a neutron star, and thus these systems are of no astrophysical interest for GW astronomy. There is no need to produce SEOBNRv2 waveforms for thsese configurations at such low frequency, but since the input waveforms need to fill the entire geometric frequency domain of interest, some data is required there. I generate the waveforms at a fixed mass of the smaller body of $1M_\odot$ and then hybridize them with the TaylorF2 stationary phase approximant. The hybridization procedure is described in Appendix A of [24].

At the high frequency end, the variation in ringdown frequency of SEOBNRv2 causes the ringdown for some configurations to start at a frequency where the ringdown has already stopped for others in the Fourier transformed data. To prevent problems with extrapolation I fit the ringdown in the frequency domain for all input waveforms, as described in Sec. IV B.

To obtain an overall ROM the low and high frequency ROMs can be glued together in frequency as follows. The phasing is glued over an overlapping frequency interval with the help of least squares fits to be C^1. This is discussed in Sec. IV A. For the amplitude I have found it sufficient to simply connect the two discrete amplitude functions at a chosen frequency in the frequency interval where they overlap. On each frequency subdomain the discrete amplitude and phase functions are given by evaluating $\mathcal{B}\cdot I_\otimes[\mathcal{M}](\lambda)$ for a specified point $\lambda=\{\eta,\chi_1,\chi_2\}$ in the parameter space.

IV. ANATOMY OF THE SEOBNRV2 ROM

Here I discuss in detail the subdomains in frequency and model parameters η,χ_1,χ_2 from which the overall ROM is constructed. I use a single domain covering the entire parameter space of SEOBNRv2 at low frequency. The parameter space is split into two subregions at high frequency, with a transition at the spin of the larger BH $\chi_1=0.41$. Both subregions have some overlap around this χ_1 value. The final ROM is assembled from these three sub models.

The gluing of ROMs in frequency is discussed in Sec. IV A, the treatment of the ringdown in Sec. IV B and the ROMs on the low and high frequency subdomains are described in Sec. IV C and Sec. IV D.

A. Gluing the ROMs in frequency

I choose a transition frequency of $Mf_m=0.01$ where the low and high frequency ROMs will be connected so that waveforms generated at a total mass of $100M_\odot$ are of sufficient length. The rationale for patching in frequency has been discussed in Sec. III B.

The phasing of the low frequency ROM is first interpolated onto the grid of the high frequency ROM. A cubic polynomial is then fitted to the phasing of the low and high frequency ROM in a frequency interval $Mf\in[0.0074,0.014]$ covered by 15 sparse frequency points below and above the gluing frequency. I then compute

$$\Delta\omega_m=\left.\frac{d\phi_{\rm high}}{df}\right|_{Mf_m}-\left.\frac{d\phi_{\rm low}}{df}\right|_{Mf_m} \tag{4.1}$$

$$\Delta\phi_m=\phi_{\rm high}(Mf_m)-\phi_{\rm low}(Mf_m)-\Delta\omega_m\,Mf_m \tag{4.2}$$

$$\tag{4.3}$$

and correct the phasing of the high frequency ROM by

$$\phi_{\rm high}(f)-\Delta\omega_m\,f-\Delta\phi_m. \tag{4.4}$$

$\phi_{\rm low}(f)$ is used unchanged up to Mf_m.

B. Extending the ringdown

The ROM construction in geometric frequency requires the waveform data to be given on a common frequency grid. As Fig. 3 shows, the ringdown for highly aligned waveforms starts at a frequency where the ringdown in the FFT data of highly anti-aligned waveforms has already terminated. One solution is to extend the ringdown data from the FFT of SEOBNRv2 by a fit and use this as *input* data for building the ROM. In particular, I fit the ringdown amplitude of all input waveforms to a Lorentzian function [19]

$$\mathcal{L}(Mf)=\gamma_1\frac{(\gamma_3 Mf_{\rm DM})^2}{(Mf-Mf_{\rm RD})^2+(\gamma_3 Mf_{\rm DM})^2}e^{\frac{-\gamma_2(Mf-Mf_{\rm RD})}{\gamma_3 Mf_{\rm DM}}}, \tag{4.5}$$

The ringdown frequency $Mf_{\rm RD}$ and damping frequency $Mf_{\rm DM}$ are computed from the LAL code [52] for SEOBNRv2. γ_1 is set to the amplitude at the ringdown frequency $Mf_{\rm RD}$ and I fit γ_2 and γ_3 in the frequency interval $Mf\in[Mf_{\rm RD},0.18]$. I use the fits to extend the ringdown amplitude from the ringdown frequency $Mf_{\rm RD}$ until $Mf=0.3$ so that the ringdown part is represented for all configurations in the domain of definition of SEOBNRv2. Linear extrapolation in the phase is sufficient due to the rapid decrease in the amplitude. Since the extrapolation of the ringdown is applied to all SEOBNRv2 input waveforms, the ROM automatically incorporates this information and can be evaluated up to $Mf=0.3$.

C. Low frequency ROM

I build a low frequency ROM that covers the entire SEOBNRv2 parameter space $0.01\leq\eta\leq0.25$, $-1\leq\chi_i\leq0.99$.

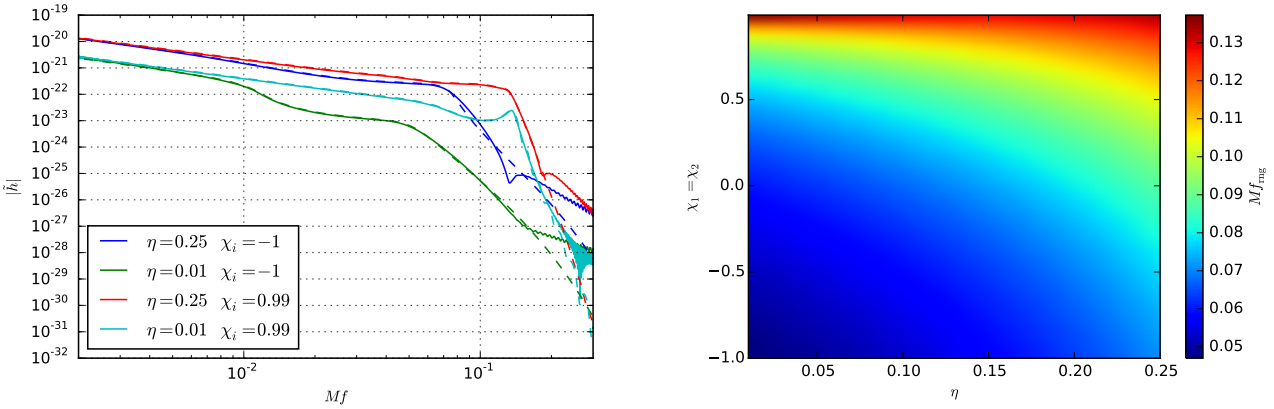

FIG. 3: The Fourier domain amplitude for four corner cases in the parameter space (left panel). The FFT of SEOBNRv2 (solid) and the ROM (dashed) are shown. The SEOBNRv2 ringdown for extreme anti-aligned configurations $\chi_i = -1$ terminates earlier than the beginning of the ringdown for $chi_i = 0.99$. The ringdown frequency of SEOBNRv2 as a function of the symmetric mass-ratio η and equal spin $\chi_1 = \chi_2$ (right panel).

The waveforms are generated at a sampling rate of 32768 Hz with a length sufficient for systems of a total mass of $2M_\odot$ or larger. Hybridization with TaylorF2 is used for the astrophysically uninteresting low frequency part for higher mass-ratios.

The set of input waveforms was defined on a $67 \times 12 \times 12$ grid in η, χ_1, χ_2 (see Fig. 4). The choice of points was obtained iteratively after checking the accuracy of the resulting ROM against SEOBNRv2 and refining the grid where needed. The sparse frequency grid used in the ROM covers $Mf \in [0.0000985, 0.15]$ using 200 points for the amplitude and 250 for the phase.

It should be pointed out that these waveforms are *very* expensive to produce. The generation time is about 1.5 hours per waveform on an Intel Xeon Haswell with 2.5GHz clock speed. Depending on the configuration the memory required per waveform can reach up to 18 GB towards the end of the EOB evolution. The $\sim$ 10000 waveforms were generated on an MPI cluster with 4 GB per core by filling the nodes only up to a third. The total cost was quite reasonable at about 50000 core hours.

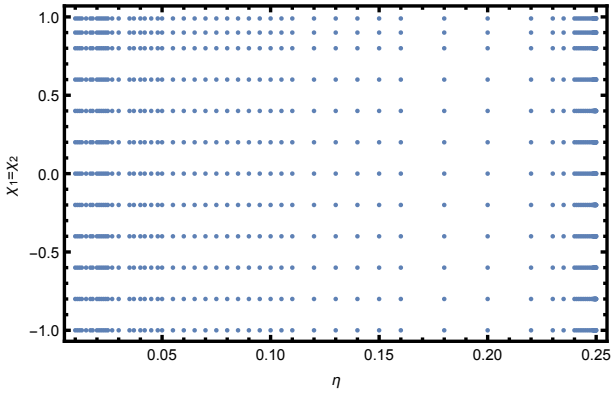

FIG. 4: The location of input waveforms in η and one of the aligned spins for the low frequency ROM. Both spin components use the same grid.

D. High frequency ROMs

At high frequency I use domain decomposition in the spin of the larger BH. The resolution requirements turned out to be more demanding for high χ_1 (see Sec. V C) and I therefore use a much finer grid for $\chi_1 \geq 0.4$ than for lower values of the spin. I describe both high frequency ROMs below, focusing on the high spin region.

Tests with a ROM covering the entire SEOBNRv2 parameter space showed unresolved structure for high spin on the larger BH at high total mass. To make a better model of this region it is efficient to work with a dense grid of short waveforms encompassing the late inspiral and merger ringdown. This model can then be combined with the low frequency ROM described in Sec. IV C.

After some refinements I chose a grid of $106 \times 119 \times 51 = 643314$ waveforms in (η, χ_1, χ_2) covering $\eta \in [0.01, 0.25]$ with $\Delta\eta \approx 0.0025$, $\chi_1 \in [0.4, 0.98999]$ with $\Delta\chi_1 \approx 0.005$, and $\chi_2 \in [-1, 0.98999]$ with $\Delta\chi_2 \approx 0.04$. For the spins, the spacings were adjusted at the boundary, to conform to the range where SEOBNRv2 waveforms could be generated reliably. (The actual allowed range is $[-1, 0.99]$, but the LAL code [52] sometimes failed at these boundaries.) To improve accuracy near the η boundaries several more points were added. The orange points in Fig. 5 show the grid in η and the spin on the larger BH, chi_1.

For $\chi_1 < 0.4$ a lower density in the spin χ_1 turned out to be sufficient for good accuracy. Fine spacing near mass-ratio boundaries was still needed. I used a $63 \times 37 \times 51$ grid in η, χ_1, χ_2. The grid in χ_2 is the same for the low and high χ_1 ROM. The χ_1 and χ_2 grids are identical up to $\chi_1 = 0.4$. The blue points in Fig. 5 show the grid in η and chi_1 for this region.

The SEOBNRv2 waveforms were generated at a total mass of $100M_\odot$ and were used in the frequency range $Mf \in$

[0.0074, 0.3] with 113 sparse frequency points for the amplitude and phase. Compared to the input waveforms for the low frequency ROM (see Sec. IV C) the computational cost for producing these high mass waveforms was small.

V. SPEEDUP AND ACCURACY

I present results on the performance of the overall ROM as defined in Sec. IV in terms of the low frequency ROM and the two ROMs at high frequency. The main criteria for a successful ROM are that it facilities data analysis applications that were infeasible with the original model and that it is accurate in terms of the match or "faithfulness" between the ROM and the original model.

The speedup relative to SEOBNRv2 is summarized in Sec. V A and the faithfulness mismatch defined in Sec. V B. Non-smooth features in SEOBNRv2 at high frequency are discussed in V C. Finally, I discuss the faithfulness of the ROM against SEOBNRv2 in Sec.V D for second and third generation ground-based GW detectors and a wide range of total masses.

A. Speedup

The speedup of the overall ROM as implemented in LAL [52] against SEOBNRv2 is shown in Fig. 6. The full number of frequency points obtained from the FFT of a time-domain waveform sampled at a rate of 16384 Hz is used. The speedup depends both on the total mass and the mass-ratio. The figure only shows the speedup down to a total mass of $12M_\odot$; for lower total mass the generation of high mass-ratio SEOBNRv2 waveforms becomes prohibitively expensive. As discussed in [24] the evaluation speed of the ROM is composed of a constant cost per frequency point for the spline interpolation in frequency and a startup cost for performing the interpolation of coefficients over the parameter space. The former dominates for long waveforms while the latter only becomes significant for very short waveforms. The speedup peaks around $20 - 50M_\odot$ reaching a factor of several thousands. For high total mass it falls off to several hundreds.

Due to gluing the ROM in frequency parameter space interpolation needs to be carried out in both frequency patches which leads to about a factor of two slowdown compared to a single-patch in frequency ROM. The speedup does not significantly depend on whether the configuration is in the low or high χ_1 patch, even though the data set is much larger for high spin. This is due to the local support of the cubic splines used to interpolate the amplitude and phase coefficients. The larger size of the high χ_1 dataset only impacts the initial step where the ROM data is loaded into memory. Since this happens only once for each analysis it does not affect performance.

Further speedup can be obtained by truncating of the expansion in amplitude and phase used in the ROM to $\sim$ 50 SVD modes (see Eq. (6.3) in [24]). This yields an additional speedup factor of 3 for high total mass and also reduces the size of the data set. This increase in speed comes at the cost of higher mismatch errors near the high mass-ratio and extreme anti-aligned spin boundaries that can reach up to 0.3%.

Some data analysis applications require a significantly smaller number of frequency points, e.g. the computation of the match in zero noise with a detector noise curve, and thus the speedup can be *much* larger than the conservative numbers shown in Fig. 6. For example parameter estimation using a match-based likelihood [53] and the generation of template banks with stochastic algorithms [6, 7] can use fairly coarse frequency spacing $O(Hz)$ without compromising the accuracy of the match computation.

Similarly, the ROM can be exploited for the construction of reduced order quadratures [54, 55] to provide further speedup in calculating inner products for low mass systems.

B. Inner product and unfaithfulness

A natural inner product between two waveforms h_1, h_2 can be defined as

$$\langle h_1, h_2\rangle := 4\mathrm{Re}\int_{f_{\rm low}}^{f_{\rm high}} \frac{\tilde{h}_1(f)\tilde{h}_2^*(f)}{S_n(f)}df, \tag{5.1}$$

where $\tilde{}$ denotes Fourier transformation, $S_n(f)$ is the power spectral density (PSD) of the detector noise and $f_{\rm low}, f_{\rm high}$ are suitable cutoff frequencies for the detector sensitivity. The lower cutoff depends on the PSD I choose below while I keep the higher cutoff fixed at 8000 Hz. The mismatch (or "unfaithfulness") is defined as the normalized inner product (5.1) maximized over time and phase shifts

$$\mathcal{M} := 1 - \max_{t_0,\phi_0} \frac{\langle h_{22}^{\rm SEOBNRv2}, h_{22}^{\rm ROM}\rangle}{||h_{22}^{\rm SEOBNRv2}||\,||h_{22}^{\rm ROM}||}, \tag{5.2}$$

where $||h|| := \sqrt{\langle h, h\rangle}$.

C. Non-smoothness in SEOBNRv2 at for high spin

Before discussing the general accuracy of the ROM in terms of unfaithfulness in Sec. V D, I first focus on the high frequency region where the detectors are sensitive to the merger-ringdown of BH binary coalescences and point out an unexpected feature in SEOBNRv2.

I compute the faithfulness against 8000 SEOBNRv2 test waveforms at a total mass of $150M_\odot$ for aLIGO design with 10 Hz starting frequency. The mismatches are smaller than 0.1% everywhere except along specific lines in the parameter space.

Fig. 7 shows a contour plot of the "comb size" parameter over which the inspiral-plunge and merger-ringdown waveforms are matched in SEOBNRv2 [22, 56] as a function of symmetric mass-ratio η and the spin of the deformed background Kerr BH χ (see Eqn. (2.1)). The comb size has a discontinuity at $\chi = 0.8$ and $\eta = 10/121$ and along these lines the mismatch of test waveforms turns out to be worse than 0.1%. These configurations are shown as red dots in Fig. 7.

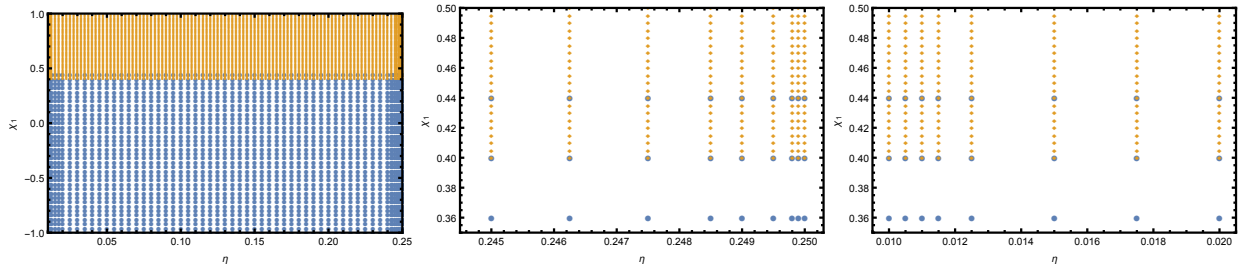

FIG. 5: The grid of input waveforms in the symmetric mass-ratio η and the aligned spin on the larger BH χ_1 for the high frequency ROMs. The points for the ROM on the low χ_1 patch are shown in blue and those for the ROM on the high χ_1 patch in orange. The left panel shows the full domain, while the middle and right panel show the transition between low and high χ_1 grids near equal mass and near the extreme mass-ratio boundary $\eta = 0.01$ of the domain.

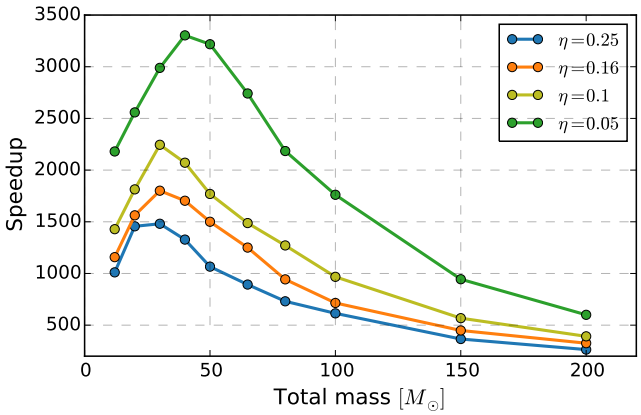

FIG. 6: Speedup of waveform generation of the overall ROM compared to SEOBNRv2.

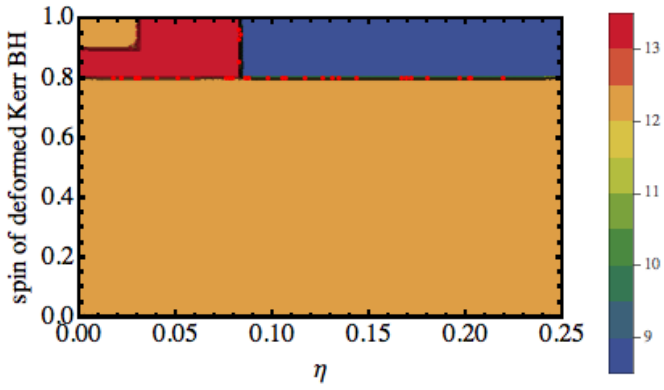

FIG. 7: The "comb size" parameter used in SEOBNRv2 is discontinuous over the parameter space as indicated by the contours. Superimposed are configurations with mismatches > 0.1% as shown in 8. They lie exactly on lines $\chi = 0.8$ and $\eta = 10/121 \approx 0.083$ where the comb size has a discontinuity.

In terms of the spin on the large BH χ_1 the discontinuity appears as a line for high mass-ratios, but for smaller mass-ratios where χ_2 becomes more important, there is significant scatter. This is shown in Fig. 8.

Non-smooth features such as the above would preclude spectral interpolation from working efficiently. Spline interpolation is more robust and the effect of non-smoothness is confined due to localized support of the basis functions.

D. Unfaithfulness against SEOBNRv2

Here I discuss the accuracy of thr ROM in terms of the faithfulness mismatch (5.2) between the ROM and SEOBNRv2, i.e. the overlap is only optimized over time and phase while mass-ratio and spins are fixed. The results were obtained from about 10000 random configurations with parameters η, χ_1, χ_2 uniformly distributed in the allowed domain. The results for aLIGO are presented in Fig. 9 for the design PSD and Fig. 10 for the early aLIGO PSD and in table I where additional data points are given as well.

Fig. 9 shows mismatches (in percent) for the aLIGO design (zero-detuned high power) PSD [48] with a lower frequency cutoff of 10 Hz. In the first three panels the mass of the smaller

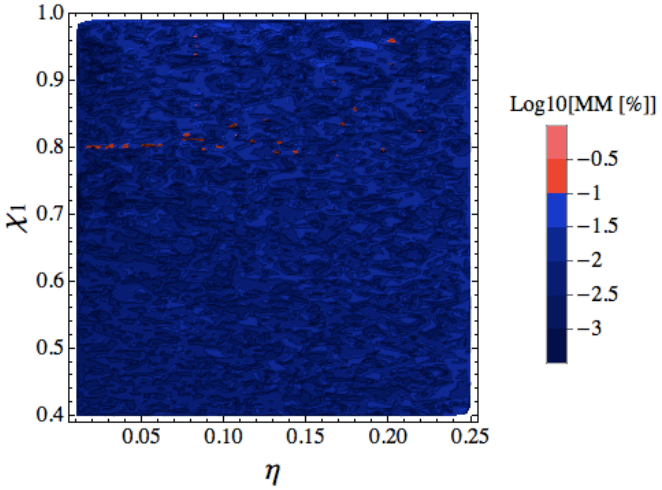

FIG. 8: Faithfulness mismatches (MM) over the parameter space as a function of the spin of the larger BH χ_1 and the symmetric mass-ratio η and as a function of the spin of the deformed background Kerr BH χ used in the EOB model.

body m_2 is fixed while the symmetric mass-ratio covers the entire domain $0.01 \leq \eta \leq 0.25$. This is done to restrict to relevant astrophysical systems with a smallest component mass above $1M_\odot$. For $m_2 = 1M_\odot$ the mismatch is in general around $\sim 0.5\%$ or lower, except for some configurations at very high mass-ratios where the mismatch can reach up to $\sim 1.5\%$. The mismatches improve considerably if the smallest component mass is increased to be $m_2 = 1.2M_\odot$ and $3M_\odot$. There, the median of the mismatch distribution is $\sim 0.003\%$ and the worst mismatch 0.25%.

In the last three panels of Fig. 9 the total mass is kept constant. For $M_{\rm tot} = 20M_\odot$ the model only covers systems up to mass-ratio $q = 19$ or $\eta = 0.0475$. Beyond this value no mismatches are computed so that the mass of the light body is at least $1M_\odot$. The median mismatch is below 0.005% for $M_{\rm tot} = 20M_\odot$ and rises to 0.02% for $M_{\rm tot} = 300M_\odot$. The worst mismatch also rises with the total mass as the merger ringdown is shifted to lower frequencies in the detector sensitivity curve. At $M_{\rm tot} = 20M_\odot$ the worst mismatch is below 0.06%. It rises to 0.4% and 1.5% for 100 and $300M_\odot$. As can also be seen from the histogram in the left panel of Fig. 11 the mismatch is higher than the continuum for a few configurations only. These configurations lie close to the lines where the SEOBNRv2 comb size parameter has a discontinuity and the waveform does not depend smoothly on parameters (see Sec. V C). After removing these points, the worst mismatch at high total mass is $\sim 0.3\%$.

Analogous to the results for aLIGO design in Fig. 9 I show mismatches for the early aLIGO PSD [57] with a lower frequency cutoff of 30 Hz in Fig. 10. The accuracy at low masses is even better for the early aLIGO noise curve. The median mismatch is below 0.003% and the worst mismatch is below 0.1% except at very high mass-ratio where it can reach 0.4% for $M_{\rm tot} = 3M_\odot$. As the total mass increases the median mismatch goes from 0.003% at $20M_\odot$ to 0.04% at $300M_\odot$. The worst mismatch rises from 0.03% to 1% and 1.9% at $20, 100, 300M_\odot$. The worst mismatches are again due to the non-smoothness of SEOBNRv2. After removing the points very close to where the comb size parameter is discontinous the worst mismatch becomes 0.2% and 0.8% at $100, 300M_\odot$. Compared to the results for aLIGO design there is some degradation of the accuracy for very high total masses. This is mostly due to the different shape of the early aLIGO PSD which tends to highlight waveform disagreement at high frequencies.

Finally, I compute mismatches using the ET-D noise-curve [58] for the third generation Einstein Telescope [4, 59]. Since the ROM is defined on a grid in geometric frequency Mf, the low frequency cutoff can be decreased by simultaneously increasing the total mass by the same factor. It is then possible to use the ROM for systems with a lighter component larger than $2M_\odot$ at a 5 Hz lower cutoff (see Fig. 12) or $5M_\odot$ at a 2 Hz lower cutoff (see Fig. 13). In the left panel of Fig. 12 the component mass was increased to $2.5M_\odot$. Then the worst mismatch is below $\sim 0.6\%$ with the median at 0.005%. The low mass accuracy in the left panel of Fig. 13 with the 2 Hz lower cutoff is comparable. The accuracy at high mass is also comparable between the two choices of lower frequency cutoff. The worst mismatch is below $\sim 0.3\%$ with the median at 0.01%.

VI. CONCLUSIONS

I have presented a reduced order (or surrogate) model of SEOBNRv2 [22] for the dominant harmonic of the GW emitted by BH binary systems. The ROM covers the entire domain of definition of SEOBNRv2: $0.01 \leq \eta \leq 0.25$ and $-1 \leq \chi_i \leq 0.99$. The ROM can be used for compact binaries from a total mass of $2M_\odot$ and larger with a 10 Hz lower frequency cutoff, or more generally for geometric frequencies from $Mf_{\rm low} = 2M_\odot[s]10{\rm Hz} = 0.0000985$ to $Mf_{\rm high} = 0.3$. This covers the entire mass range range of BNS, NS-BH and BBH sytems of interest for second generation ground-based laser-interferometric detectors aLIGO and AdvVIRGO. The ROM can also be use for data analysis with ET for slightly heavier systems from $\sim 2 - 5M_\odot$ onwards depending on the lower frequency cutoff.

While this ROM is based on the method described in [24], the construction relies on extensions to these techniques (see Secs. III and IV). In particular, the parameter space is decomposed into two 3-dimensional patches in η, χ_1, χ_2 and a division into a low and a high frequency patch in Mf. This domain decomposition technique makes it possible to model the smooth inspiral part with a coarser waveform set where waveform generation is most costly and to obtain higher resolution for the late inspiral and merger-ringdown, where waveform generation is much cheaper. The ringdown is fit against a Lorentzian to guarantee that the complete ringdown is modeled for all configurations in the SEOBNRv2 parameter space.

As shown in Fig. 6 this ROM is several thousand times faster than SEOBNRv2 when using the full number of frequency points. This speedup is crucial for searches and parameter estimation in the LSC [60]. The SEOBNRv2 ROM has already been used in theoretical parameter estimations studies within the LSC [61, 62] and is used in aLIGO's first observing run. Data analysis applications such as the generation of stochastic template banks for searches [6–10] and parameter estimation studies in zero noise [53] that can elucidate important properties of waveform models require only the computation of matches with a PSD. For these applications speedups can be several tens of thousands.

Very recently an optimized version of the SEOBNRv2 code has become available [63] in LALSuite [52], although it has not yet been reviewed in the LSC. This code provides a speedup of one to two orders of magnitude over the reviewed SEOBNRv2, with the largest speedups observed for the longest waveforms (e.g., binary neutron stars). However, the ROM presented here is still more than an order of magnitude faster than this optimized code if all frequency points are used and can be a factor thousand faster for specific data analysis applications mentioned above. The availability of optimized EOB codes is very beneficial for the construction of ROMs.

A previous version of this ROM (`SEOBNRv2_ROM_DoubleSpin`) has been reviewed by

Unfaithfulness	$m_2 = 1M_\odot$	$1.2M_\odot$	$3M_\odot$	$M = 20M_\odot$	$50M_\odot$	$80M_\odot$	$100M_\odot$	$200M_\odot$	$300M_\odot$
aLIGO design 10 Hz									
$max_\lambda \mathcal{M}$[%]	1.5 (0.03 for $\eta > 0.05$)	0.22	0.27	0.06	0.19	0.53 (0.3)	0.45 (0.3)	1.0 (0.26)	1.5 (0.3)
$med_\lambda \mathcal{M}$[%]	0.02	0.002	0.003	0.004	0.009	0.14	0.016	0.02	0.024
aLIGO design 20 Hz									
$max_\lambda \mathcal{M}$[%]	0.23	0.11	0.29	0.03	0.18	0.34 (0.2)	0.47 (0.2)	0.9 (0.25)	1.37 (0.3)
$med_\lambda \mathcal{M}$[%]	0.001	0.001	0.003	0.004	0.009	0.14	0.016	0.02	0.024
aLIGO early 30 Hz									
$max_\lambda \mathcal{M}$[%]	0.05	0.08	0.38	0.026	0.15	0.63 (0.13)	1.1 (0.18)	1.6 (0.4)	1.9 (0.76)
$med_\lambda \mathcal{M}$[%]	0.0006	0.0007	0.003	0.003	0.01	0.015	0.018	0.029	0.04

TABLE I: Faithfulness mismatches of the ROM against SEOBNRv2 for alIGO design with a lower cutoff frequency of 10 or 20 Hz and early aLIGO with 30 Hz. The mismatch is given for a range of fixed masses of the smaller body m_2 and for a range of fixed total masses M. I quote both the maximum and the median of the mismatch distribution in %. For high total masses $M \geq 80M_\odot$ I give the overall maximum mismatch and the maximum mismatch with points near $\chi \sim 0.8$ removed where SEOBNRv2 is non-smooth in parentheses.

the LSC and is used to build template banks for searches and as a model in parameter estimation simulations during the first observing run of aLIGO. An implementation of the ROM described in this paper is available under the name of `SEOBNRv2_ROM_DoubleSpin_HI` in the LSC Algorithmic Library Suite (LALSuite) [52]. Compared to the previous version the ROM presented in this paper better resolves the region where SEOBNRv2 is non-smooth, introduces patching in frequency, and thereby significantly extends the total mass range from $12M_\odot$ down to BNS systems with $2M_\odot$.

Apart from the speedup, a ROM must provide a good approximation of the original waveform. This can be quantified by the match or the faithfulness between the two models over the parameter space. The unfaithfulness for SEOBNRv2 against SXS waveforms was found to be less than 1% for total masses from $20M_\odot$ to $200M_\odot$ [22]. In this mass range the mismatch between the ROM and SEOBNRv2 is always well below the SEOBNRv2 calibration accuracy except along lines where SEOBNRv2 has been shown to be non-smooth (see Sec. V C) for high total masses. Below $20M_\odot$ the ROM agrees in general extremely well (mismatch $\lesssim 0.1\%$) with SEOBNRv2. Only for extreme NS-BH systems with very high mass-ratio and a very light NS with a 10 Hz lower cutoff does the mismatch reach 1%.

Beyond faithfulness GW searches and PE require the location of the amplitude maximum in the time domain strain is accurate. A describes how a time shift can be applied to the ROM phasing in order to achieve this. Moreover, useful functions can be defined for the length of a waveform in time as a function of a starting frequency and its inverse.

The techniques described here can be applied to building ROMs for aligned-spin waveform models including higher harmonics, with a mode-by-mode approach [64]. Modeling GWs emitted from precessing binaries is a much harder problem to tackle due to the significantly larger size of the spin space and will require the development of further techniques.

Appendix A: Chirp Time and correction of time of coalescence

A "time of frequency" function and its inverse can be defined from the Fourier domain phasing. These functions are useful when the template duration or "chirp time" [65] is needed by data analysis algorithms and also allow to correct the phasing of the ROM such that the location of the amplitude peak in the time domain agrees with SEOBNRv2 to good accuracy.

One can define the notion of *time* from the frequency domain phasing by analogy with *frequency* in the time domain (in a stationary phase sense) as

$$t_{\rm FD}(Mf)/M = \frac{1}{2\pi}\frac{d\phi}{d(Mf)}. \tag{A1}$$

At low frequency this function is well approximated by a power law [65] $t_{\rm FD}(Mf)/M \propto (Mf)^{-8/3}$ and flattens out as the SEOBNRv2 ringdown frequency $f_{\rm rng}$, the highest frequency in the corresponding time domain waveform, is approached. Thus I take $f_{\rm rng}$ as a natural reference point and compute the time elapsed from a starting frequency $f_{\rm start}$ until $f_{\rm rng}$ which demarcates the end of the waveform.

$$\Delta T/M = t_{\rm FD}(f_{\rm rng})/M - t_{\rm FD}(f_{\rm start})/M \tag{A2}$$

To provide the inverse function, frequency as a function of time, $Mf(t_{\rm FD})$ I generate data for $t_{\rm FD}(Mf)$ on a double-logarithmic grid, spline interpolate the transposed data (frequency as a function of time) to find $I[\log Mf](\log t_{\rm FD}/M)$ and obtain

$$Mf(t_{\rm FD}) = \exp(I[\log Mf](\log t_{\rm FD}/M)) \tag{A3}$$

In order to avoid pathologies close to merger where $t_{\rm FD}$ is not necessarily monotonically decreasing, the double logarithmic grid extends from the ROM starting frequency until half of the ringdown frequency.

In the ROM construction an implicit choice was made for the time and phase of coalescence. These parameters are maximized over when computing the match and thus have not been discussed so far in this paper. GW searches and parameter estimation expect the amplitude peak in the time domain

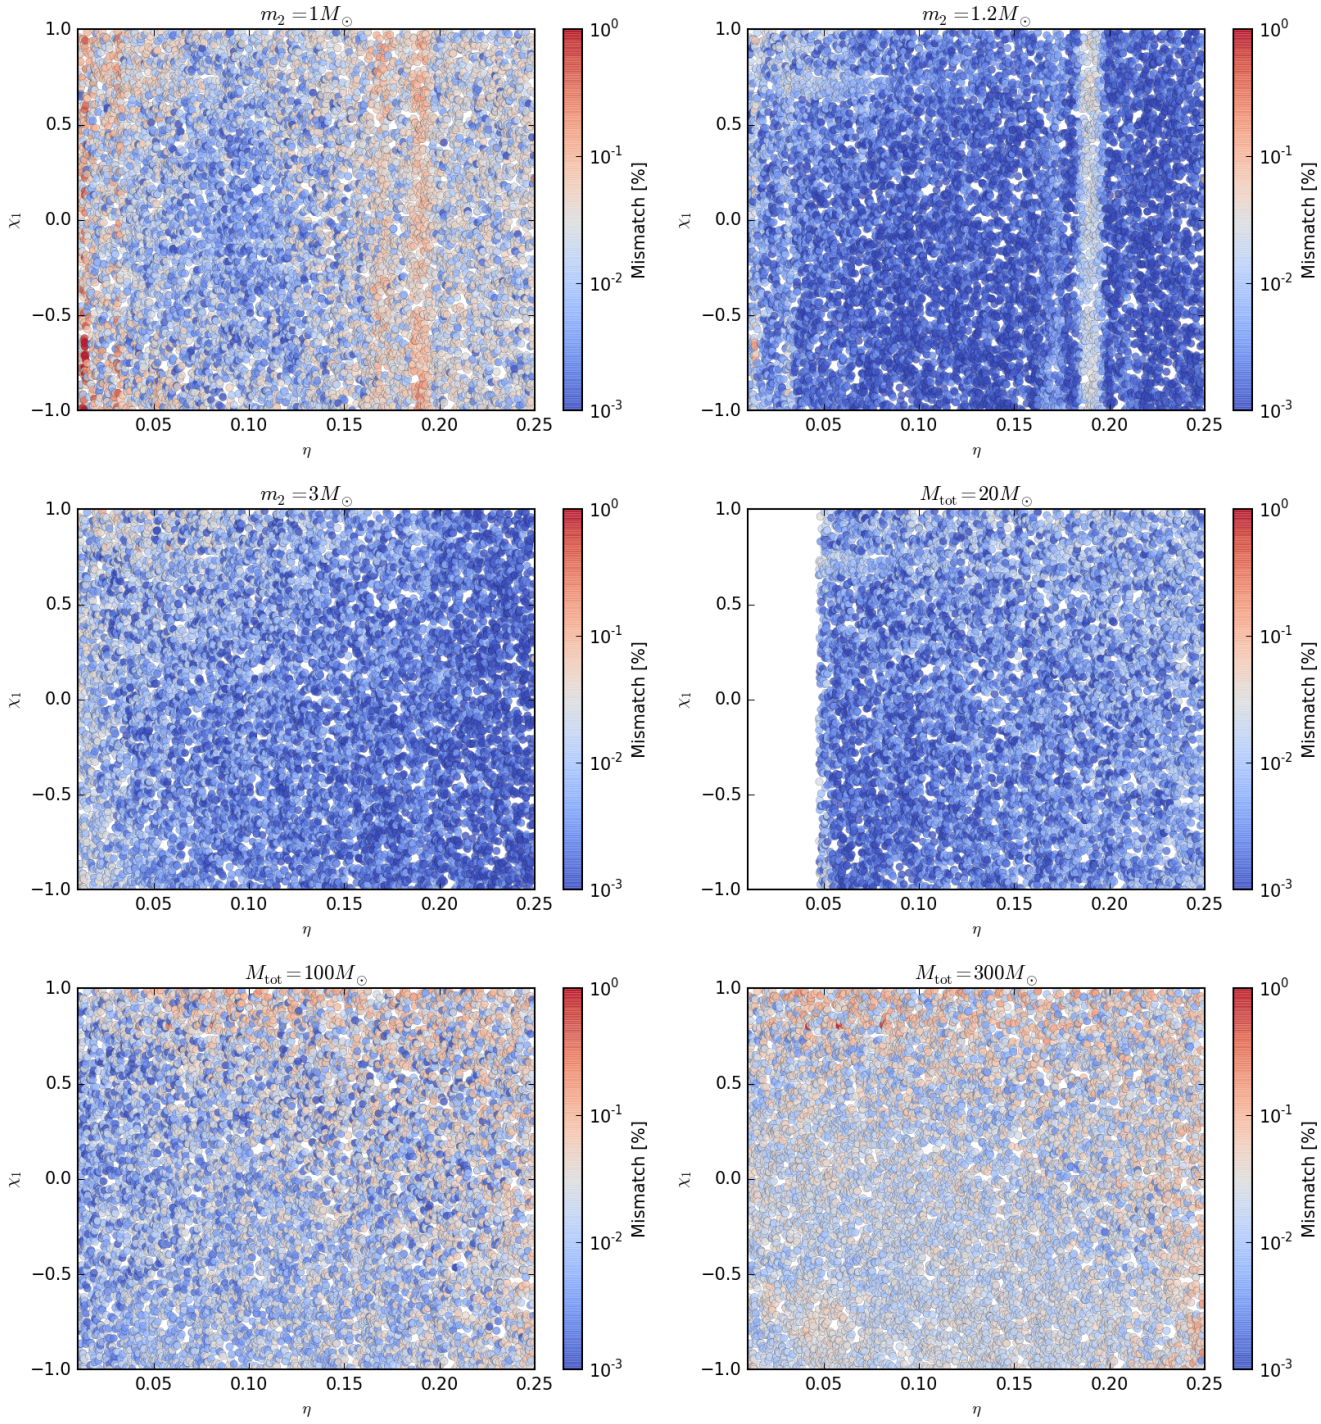

FIG. 9: Mismatches between the overall ROM and SEOBNRv2 for aLIGO design with a low frequency cutoff of 10 Hz. To facilitate comparisons between panels, the colorbar scales are fixed at a mismatch of 1%. The worst mismatches at high mass lie on lines where SEOBNRv2 is non-smooth (see Sec. V C). Histograms of the mismatches are shown in Fig. 11 (left panel).

to occur within a prescribed time window and customarily set the origin of time to the location of the peak. To ensure that the inverse Fourier transform of the ROM follows this convention I apply a time shift to the Fourier domain phasing such that $d\phi/d(Mf)|_{Mf_{\rm mg}} = 0$.

Fig. 14 shows that the maximum error in the location of the time domain amplitude peak is about 25ms at a total mass of $100M_\odot$. The ROM peak occurs up to $\sim$ 20ms too early for very high aligned spin systems and close to equal mass systems at moderate to high spin. For moderate spins and

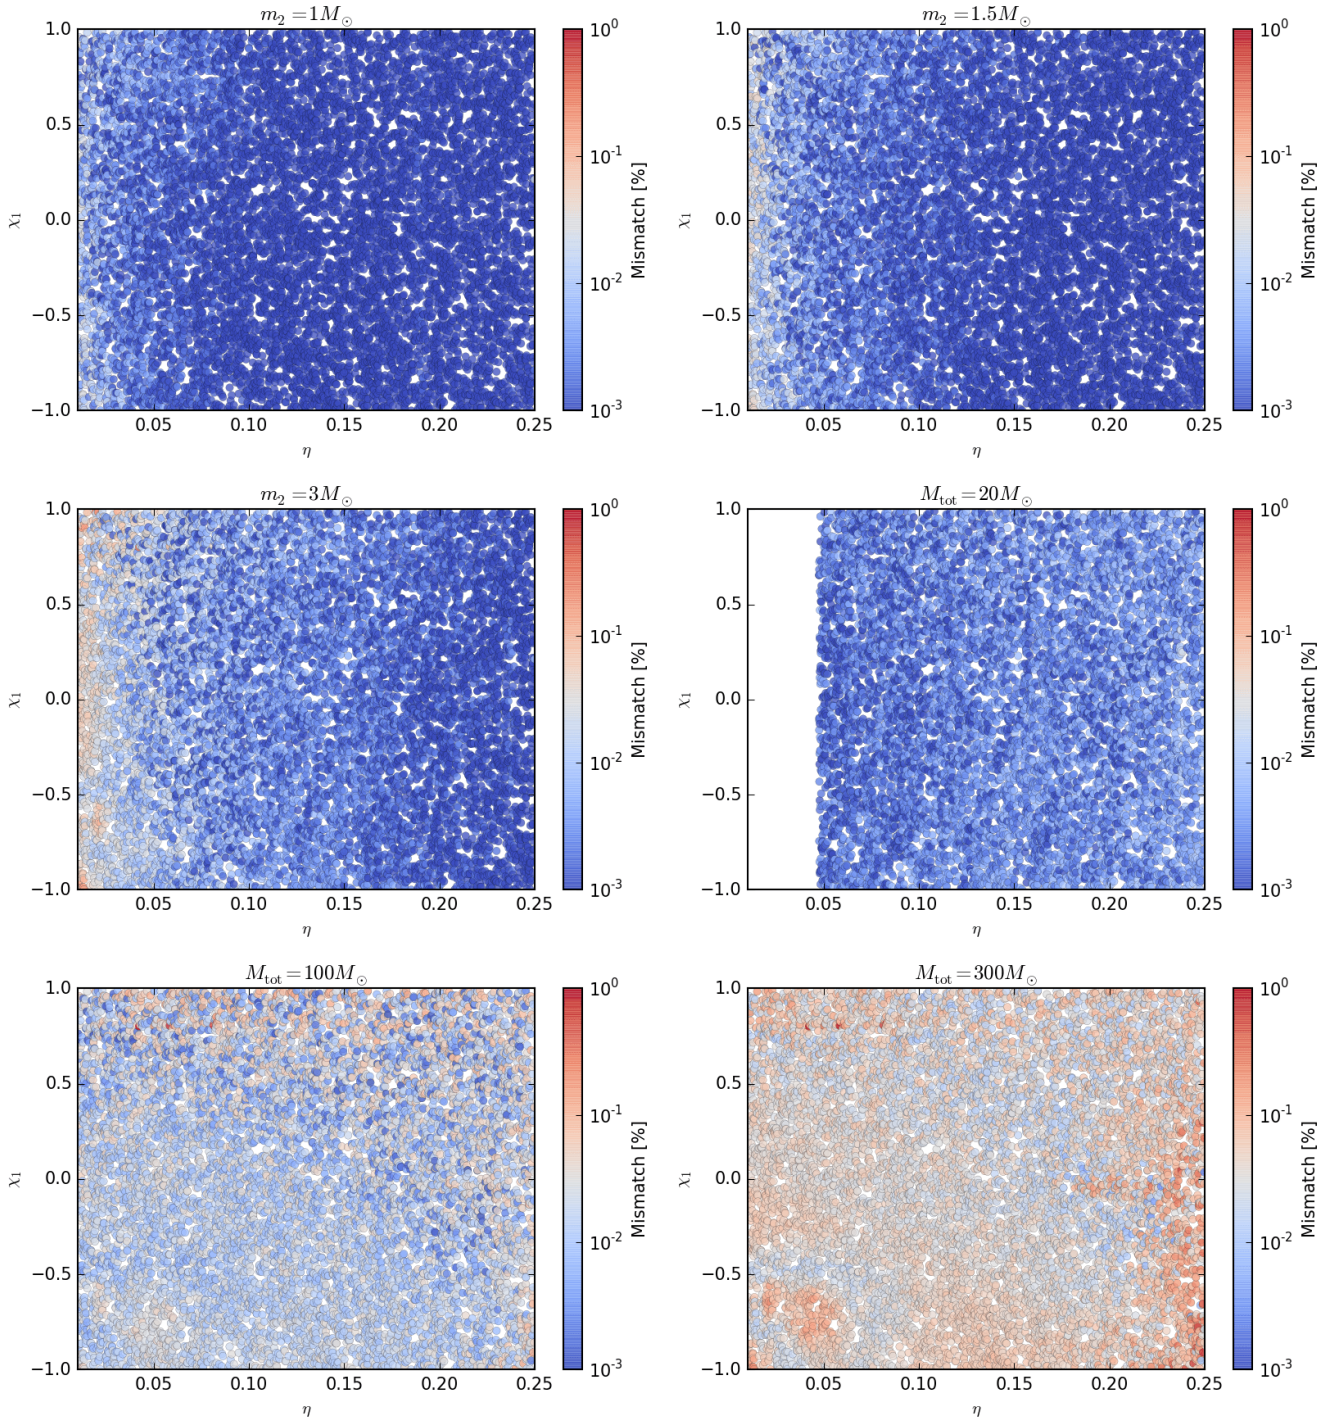

FIG. 10: Mismatches between the overall ROM and SEOBNRv2 for early aLIGO with a low frequency cutoff of 30 Hz. To facilitate comparisons between panels, the colorbar scales are fixed at a mismatch of 1%. As in Fig. 9 the worst mismatches at high mass lie on lines where SEOBNRv2 is non-smooth (see Sec. V C). Histograms of the mismatches are shown in Fig. 11 (right panel).

mass-ratios the peak offset is about −10ms. The maximum peak time error is similar for unequal spin configurations. This deviation scales linearly with the total mass and is within the time windows used by searches and parameter estimation within the LAL codes [52] used in the LSC [60].

Acknowledgements

I thank Mark Hannam, Andrea Taracchini, Alessandra Buonanno, Frank Ohme, Prayush Kumar, Harald Pfeiffer and Chad Galley for useful discussions and comments.

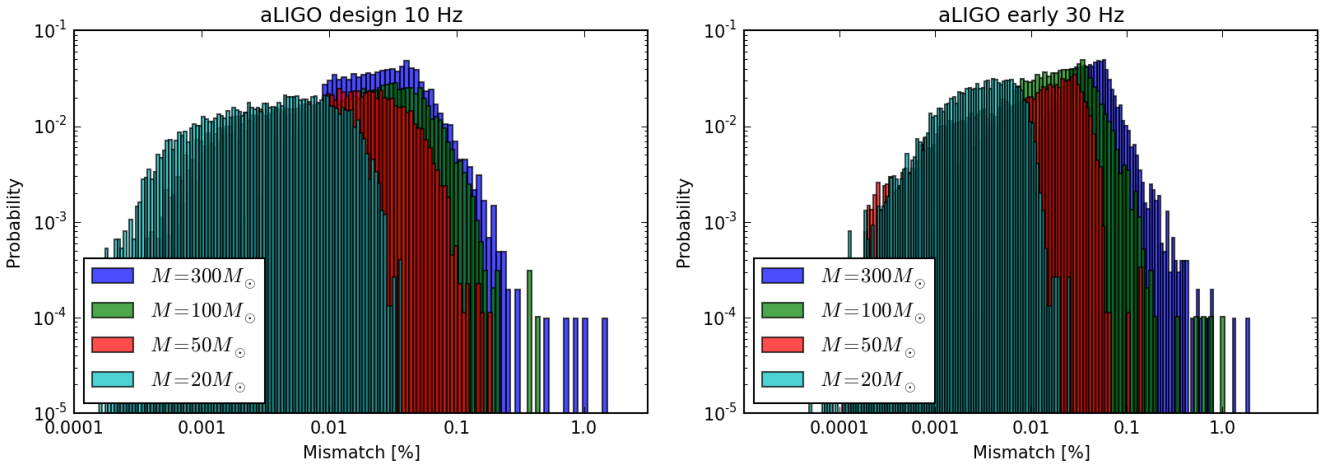

FIG. 11: Mismatch histograms for aLIGO zdethp 10 Hz (left) and aLIGO early 30 Hz (right). At high total mass isolated points have mismatch close to (or above) 1%. These configurations are due to non-smoothness in SEOBNRv2.

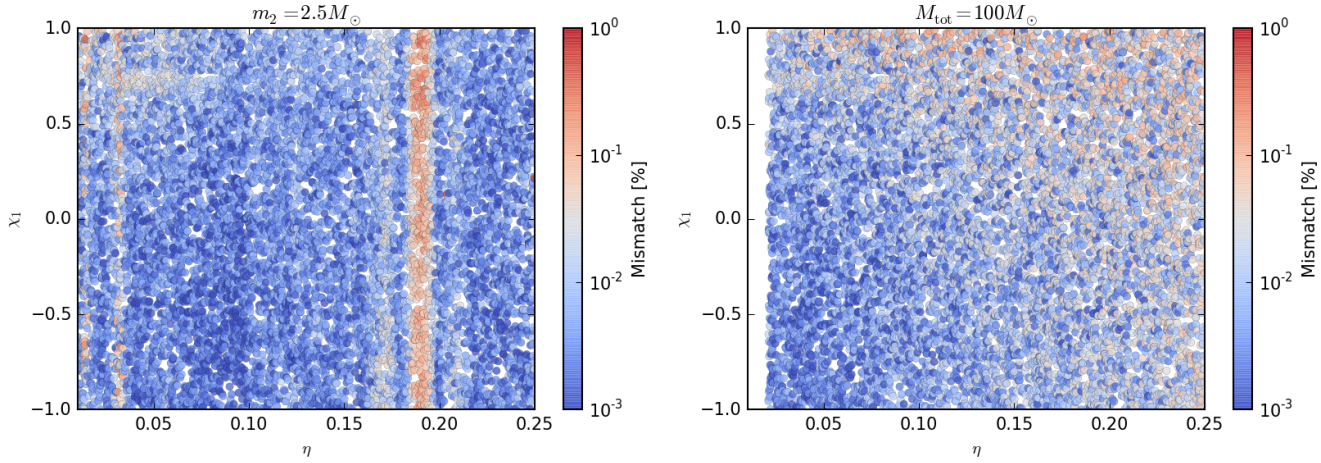

FIG. 12: Mismatches between the (glued) ROM and SEOBNRv2 for ET-D with a low frequency cutoff of 5 Hz.

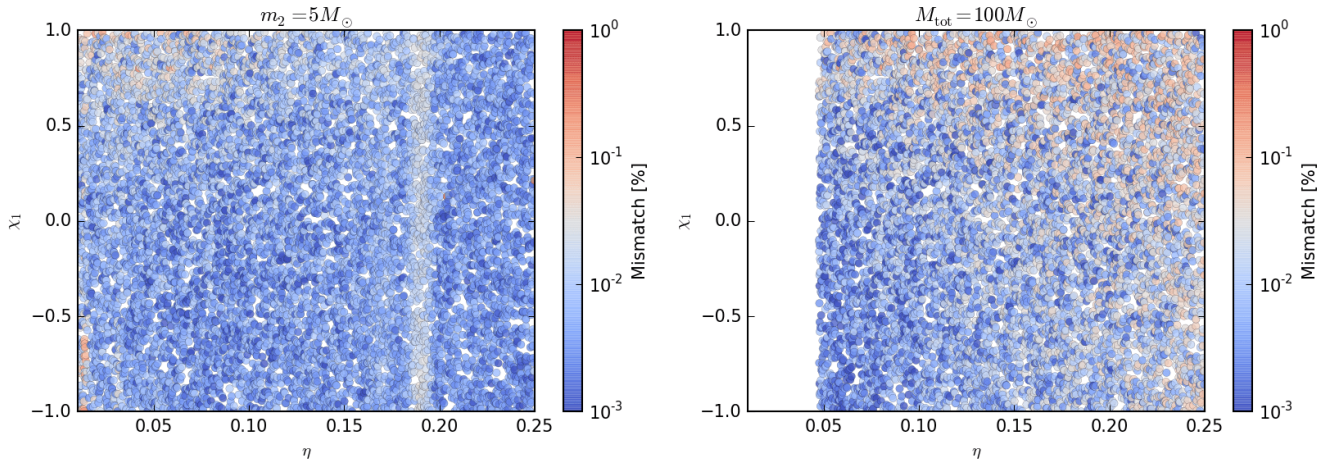

FIG. 13: Mismatches between the (glued) ROM and SEOBNRv2 for ET-D with a low frequency cutoff of 2 Hz.

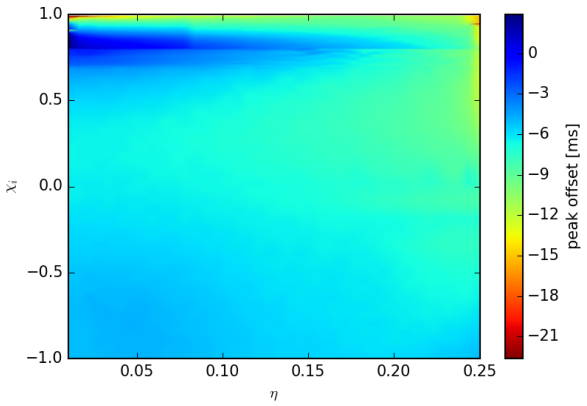

FIG. 14: Offset between the time where the amplitude in the time domain peaks for SEOBNRv2 and ROM. If the offset is negative (positive) the peak in the ROM precedes (follows) the peak in SEOBNRv2. The offset is given in milliseconds for configurations at a total mass of $100M_\odot$ as a function of the symmetric mass-ratio η and equal aligned spins χ_i.

MP was in part supported by Science and Technology Facilities Council grant ST/I001085/1. SEOBNRv2 waveforms were generated at Advanced Research Computing (ARCCA) at Cardiff.

References

[1] T. L. S. Collaboration, J. Aasi, B. P. Abbott, R. Abbott, T. Abbott, M. R. Abernathy, K. Ackley, C. Adams, T. Adams, P. Addesso, et al., Classical and Quantum Gravity **32**, 074001 (2015), URL `http://stacks.iop.org/0264-9381/32/i=7/a=074001`.

[2] F. Acernese, M. Agathos, K. Agatsuma, D. Aisa, N. Allemandou, A. Allocca, J. Amarni, P. Astone, G. Balestri, G. Ballardin, et al., Classical and Quantum Gravity **32**, 024001 (2015), URL `http://stacks.iop.org/0264-9381/32/i=2/a=024001`.

[3] J. Aasi et al. (LIGO Scientific Collaboration, Virgo Collaboration) (2013), 1304.0670.

[4] M. Punturo et al., Class. Quant. Grav. **27**, 084007 (2010).

[5] B. Allen, W. G. Anderson, P. R. Brady, D. A. Brown, and J. D. Creighton, Phys.Rev. **D85**, 122006 (2012), gr-qc/0509116.

[6] I. W. Harry, B. Allen, and B. Sathyaprakash, Phys.Rev. **D80**, 104014 (2009), 0908.2090.

[7] G. M. Manca and M. Vallisneri, Phys.Rev. **D81**, 024004 (2010), 0909.0563.

[8] P. Ajith, N. Fotopoulos, S. Privitera, A. Neunzert, N. Mazumder, and A. J. Weinstein, Phys. Rev. D **89**, 084041 (2014), URL `http://link.aps.org/doi/10.1103/PhysRevD.89.084041`.

[9] S. Privitera, S. R. P. Mohapatra, P. Ajith, K. Cannon, N. Fotopoulos, M. A. Frei, C. Hanna, A. J. Weinstein, and J. T. Whelan, Phys. Rev. D **89**, 024003 (2014), URL `http://link.aps.org/doi/10.1103/PhysRevD.89.024003`.

[10] I. W. Harry, A. H. Nitz, D. A. Brown, A. P. Lundgren, E. Ochsner, and D. Keppel, Phys. Rev. D **89**, 024010 (2014), URL `http://link.aps.org/doi/10.1103/PhysRevD.89.024010`.

[11] I. Harry, private communication (2015).

[12] S. A. Usman et al. (2015), 1508.02357.

[13] A. Buonanno and T. Damour, Phys. Rev. **D59**, 084006 (1999), gr-qc/9811091.

[14] A. Buonanno and T. Damour, Phys. Rev. **D62**, 064015 (2000).

[15] T. Damour, P. Jaranowski, and G. Schaefer, Phys.Rev. **D62**, 084011 (2000), gr-qc/0005034.

[16] T. Damour, Phys.Rev. **D64**, 124013 (2001), gr-qc/0103018.

[17] T. Damour, B. R. Iyer, and A. Nagar, Phys.Rev. **D79**, 064004 (2009), 0811.2069.

[18] S. Husa, S. Khan, M. Hannam, M. Pürrer, F. Ohme, X. J. Forteza, and A. Bohé (2015), 1508.07250.

[19] S. Khan, S. Husa, M. Hannam, F. Ohme, M. Pürrer, X. J. Forteza, and A. Bohé (2015), 1508.07253.

[20] P. Ajith et al., Phys. Rev. Lett. **106**, 241101 (2011), 0909.2867.

[21] L. Santamaria et al., Phys. Rev. **D82**, 064016 (2010), 1005.3306.

[22] A. Taracchini, A. Buonanno, Y. Pan, T. Hinderer, M. Boyle, et al., Phys.Rev. **D89**, 061502 (2014), 1311.2544.

[23] A. Taracchini, Y. Pan, A. Buonanno, E. Barausse, M. Boyle, et al., Phys.Rev. **D86**, 024011 (2012), 1202.0790.

[24] M. Pürrer, Class. Quant. Grav. **31**, 195010 (2014), 1402.4146.

[25] R. Smith, K. Cannon, C. Hanna, D. Keppel, and I. Mandel (2012), 1211.1254.

[26] K. Cannon, J. Emberson, C. Hanna, D. Keppel, and H. Pfeiffer, Phys.Rev. **D87**, 044008 (2013), 1211.7095.

[27] K. Cannon, C. Hanna, and D. Keppel, Phys.Rev. **D85**, 081504 (2012), 1108.5618.

[28] S. E. Field, C. R. Galley, J. S. Hesthaven, J. Kaye, and M. Tiglio, Phys. Rev. **X4**, 031006 (2014), 1308.3565.

[29] S. E. Field, C. R. Galley, F. Herrmann, J. S. Hesthaven, E. Ochsner, et al., Phys.Rev.Lett. **106**, 221102 (2011), 1101.3765.

[30] J. Blackman, B. Szilagyi, C. R. Galley, and M. Tiglio, Phys.Rev.Lett. **113**, 021101 (2014), 1401.7038.

[31] J. Blackman, S. E. Field, C. R. Galley, B. Szilágyi, M. A. Scheel, M. Tiglio, and D. A. Hemberger, Phys. Rev. Lett. **115**, 121102 (2015), 1502.07758.

[32] F. Pannarale, E. Berti, K. Kyutoku, B. D. Lackey, and M. Shibata, Phys. Rev. **D92**, 084050 (2015), 1509.00512.

[33] B. D. Lackey, K. Kyutoku, M. Shibata, P. R. Brady, and J. L. Friedman, Phys. Rev. **D89**, 043009 (2014), 1303.6298.

[34] M. Fares, J. Hesthaven, Y. Maday, and B. Stamm, Journal of Computational Physics **230**, 5532 (2011), ISSN 0021-9991, URL http://www.sciencedirect.com/science/article/pii/S0021999111001665.

[35] J. L. Eftang, D. J. Knezevic, and A. T. Patera, Mathematical and Computer Modelling of Dynamical Systems **17**, 395 (2011), http://dx.doi.org/10.1080/13873954.2011.547670, URL http://dx.doi.org/10.1080/13873954.2011.547670.

[36] J. Eftang, D. Huynh, D. Knezevic, and A. Patera, Journal of Scientific Computing **51**, 28 (2012), ISSN 0885-7474, URL http://dx.doi.org/10.1007/s10915-011-9494-2.

[37] J. Eftang and B. Stamm, Int. J. Numer. Methods Eng. **90**, 412 (2012).

[38] L. Blanchet, Living Rev.Rel. **17**, 2 (2014), 1310.1528.

[39] J. Centrella, J. G. Baker, B. J. Kelly, and J. R. van Meter, Rev.Mod.Phys. **82**, 3069 (2010), 1010.5260.

[40] E. Berti, V. Cardoso, and A. O. Starinets, Class. Quant. Grav. **26**, 163001 (2009), 0905.2975.

[41] A. Buonanno and B. S. Sathyaprakash (2014), 1410.7832.

[42] T. Damour and A. Nagar, Phys.Rev. **D79**, 081503 (2009), 0902.0136.

[43] Y. Pan, A. Buonanno, M. Boyle, L. T. Buchman, L. E. Kidder, et al., Phys.Rev. **D84**, 124052 (2011), 1106.1021.

[44] A. H. Mroue et al., Phys. Rev. Lett. **111**, 241104 (2013), 1304.6077.

[45] A. H. Mroue and H. P. Pfeiffer (2012), 1210.2958.

[46] D. A. Hemberger, G. Lovelace, T. J. Loredo, L. E. Kidder, M. A. Scheel, B. Szilágyi, N. W. Taylor, and S. A. Teukolsky, Phys. Rev. D **88**, 064014 (2013), URL http://link.aps.org/doi/10.1103/PhysRevD.88.064014.

[47] D. A. Hemberger, M. A. Scheel, L. E. Kidder, B. Szilágyi, G. Lovelace, N. W. Taylor, and S. A. Teukolsky, Class. Quant. Grav. **30**, 115001 (2013), 1211.6079.

[48] The LIGO Scientific Collaboration, Tech. Rep. LIGO-T0900288-v3, LIGO Project (2009), URL https://dcc.ligo.org/DocDB/0002/T0900288/003/AdvLIGO%20noise%20curves.pdf.

[49] P. Kumar, K. Barkett, S. Bhagwat, N. Afshari, D. A. Brown, G. Lovelace, M. A. Scheel, and B. Szilágyi, Phys. Rev. **D92**, 102001 (2015), 1507.00103.

[50] P. Kumar, T. Chu, H. P. Pfeiffer, M. Boyle, C. Hadmer, D. Hemberger, L. Kidder, M. A. Scheel, and B. Szilagyi, LIGO-P1500251.

[51] T. Chu, H. Fong, P. Kumar, H. P. Pfeiffer, M. Boyle, C. Hadmer, D. Hemberger, L. E. Kidder, M. A. Scheel, and B. Szilagyi, LIGO-P1500250.

[52] LALSuite, The LIGO Scientific Collaboration, Lsc algorithm library (lal) – https://www.lsc-group.phys.uwm.edu/daswg/projects/lalsuite.html.

[53] M. Pürrer, M. Hannam, and F. Ohme, in preparation.

[54] P. Canizares, S. E. Field, J. R. Gair, and M. Tiglio, Phys. Rev. **D87**, 124005 (2013), 1304.0462.

[55] P. Canizares, S. E. Field, J. Gair, V. Raymond, R. Smith, and M. Tiglio, Phys. Rev. Lett. **114**, 071104 (2015), 1404.6284.

[56] A. Taracchini, private communication (2015).

[57] Shoemaker, David and Losurdo, Giovanni, Tech. Rep. LIGO-G10001760-v7, LIGO Project (2010), URL https://dcc.ligo.org/DocDB/0009/G1000176/007/Advanced%20Virgo%20and%20LIGO%20post-project%20progression.pdf.

[58] S. Hild, M. Abernathy, F. Acernese, P. Amaro-Seoane, N. Andersson, K. Arun, F. Barone, B. Barr, M. Barsuglia, M. Beker, et al., Classical and Quantum Gravity **28**, 094013 (2011), URL http://stacks.iop.org/0264-9381/28/i=9/a=094013.

[59] B. Sathyaprakash et al., Class. Quant. Grav. **29**, 124013 (2012), [Erratum: Class. Quant. Grav.30,079501(2013)], 1206.0331.

[60] LSC, LSC – http://www.ligo.org.

[61] J. Veitch, M. Pürrer, and I. Mandel, Phys. Rev. Lett. **115**, 141101 (2015), 1503.05953.

[62] C.-J. Haster, Z. Wang, C. P. L. Berry, S. Stevenson, J. Veitch, and I. Mandel (2015), 1511.01431.

[63] C. Devine, Z. B. Etienne, and S. T. McWilliams, in preparation.

[64] S. Marsat, P. B. Graff, J. G. Baker, and A. Buonanno, in preparation.

[65] B. S. Sathyaprakash, Phys. Rev. D **50**, R7111 (1994), URL http://link.aps.org/doi/10.1103/PhysRevD.50.R7111.